Creating A Musical Performance Dataset for Multimodal Music Analysis: Challenges, Insights, and Applications

Bochen Li*, Xinzhao Liu*, Karthik Dinesh, Zhiyao Duan, *Member, IEEE,* and Gaurav Sharma, *Fellow, IEEE*

arXiv:1612.08727v1 [cs.MM] 27 Dec 2016

***Abstract*—We introduce a dataset for facilitating audio-visual analysis of musical performances. The dataset comprises a number of simple multi-instrument musical pieces assembled from coordinated but separately recorded performances of individual tracks. For each piece, we provide the musical score in MIDI format, the audio recordings of the individual tracks, the audio and video recording of the assembled mixture, and ground-truth annotation files including frame-level and note-level transcriptions. We anticipate that the dataset will be useful for developing and evaluating multi-modal techniques for music source separation, transcription, score following, and performance analysis. We describe our methodology for the creation of this dataset, particularly highlighting our approaches for addressing the challenges involved in maintaining synchronization and naturalness. We briefly discuss the research questions that can be investigated with this dataset.**

***Index Terms*—Multimodal music dataset, audio-visual analysis, music performance, synchronization.**

I. Introduction

MUSIC performance is a multimodal art form. For thousands of years, people have enjoyed music performances at live concerts through both hearing and sight. The development of recording technologies, starting with Thomas Edison's invention of the gramophone in 1877, has extended music enjoyment beyond live concerts. For a long time, the majority of music recordings were distributed through various kinds of media that carry only the audio, such as vinyl records, cassettes, CDs, and mp3 files. As such, existing research on music analysis, processing, and retrieval focused on the audio modality, while the visual component was largely forgotten.

About a decade ago, with the rapid expansion of digital storage and internet bandwidth, video streaming services like YouTube gained popularity, which again significantly influenced the way people enjoy music. Not only do the audiences want to listen to the sound, they also want to watch the performance. In 2014, music was the most searched topic on YouTube, and 38.4% of YouTube video views were from music videos [1]. The visual modality plays an important role in musical performances. Guitar players learn new songs by watching how others play online. Concert audience move their gazes to the soloist in a Jazz concert. In fact, researchers have found that the visual component is not a marginal phenomenon in music perception, but an important factor in the communication of meanings [2]. Even for prestigious classical music competitions, researchers have found that visually perceived elements of the performance, such as gesture, motion, and facial expressions of the performer, affect the judge's (experts or novice alike) evaluations, even more significantly than the sound [3].

Music Information Retrieval (MIR) research, which traditionally focused on audio and symbolic modalities (e.g., musical scores), started to pay attention to the visual modality in recent years. When available, The visual components can be very helpful for solving many traditional MIR tasks that are challenging using an audio-only approach. Zhang et al. [4] introduced a method to transcribe solo violin performances by tracking the violin strings and fingers from the visual scene. Similarly, remarkable success has been demonstrated for music transcription using visual techniques for piano [5], guitar [6] and drums [7]. The use of visual information also enables detailed analyses of many attributes of musical performances, including sound articulation, playing techniques, and expressiveness. Note that these benefits of incorporating visual information in the analysis of audio are especially pronounced for highly polyphonic, multi-instrument performances, because the visual activity of each player is usually directly observable (barring occlusions), whereas the polyphony makes it difficult to unambiguously associate audio components with each player. Dinesh et al. [8] proposed to detect play/non-play activity for each player in a string ensemble to augment multi-pitch estimation and streaming. A similar idea is applied on different instrument groups among symphony orchestras to achieve performance-score alignment [9]. Additionally, audio-visual analysis opens up new frontiers for several existing and new MIR problems. Researchers have proposed systems to analyze the fingering of guitarists [10]–[13] and pianists [14], [15], the baton trajectories of the conductors [16], the audio-visual source association in multimedia music [17], and the interaction modes between players and instruments [18].

Despite the increased recent interest, progress in jointly using audio and visual modalities for the analysis of musical performances has been rather slow. One of the main reasons, we argue, is the lack of datasets. Although music takes a large share among all kinds of multimedia data, music datasets are scarce. This is because a music dataset should contain not only music recordings but also ground-truth annotations (e.g.,

*B. Li and X. Liu made equal contribution to this paper.

X. Liu was with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627, USA. She is now with Listent American Corp. E-mail: xinzhao.liu@rochester.edu

B. Li, K. Dinesh, Z. Duan, and G. Sharma are with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627, USA. E-mails: {bochenli, karthik.dinesh, zhiyao.duan, gaurav.sharma}@rochester.edu.

note/beat/chord transcriptions, performance-score alignments) to enable supervised machine learning and the evaluation of proposed methods. Due to the temporal and polyphonic nature of music, the annotation process is very time consuming and often requires significant musical expertise. Furthermore, for some research problems such as source separation, isolated recordings of different sound sources (e.g., musical instruments) are also needed for ground truth verification. When creating such a dataset, if each source is recorded in isolation, it is a challenging task to ensure that different sources are well tuned and properly synchronized.

In this paper, we present the University of Rochester Musical Performance (URMP) dataset. This dataset covers 44 classical chamber music pieces ranging from duets to quintets. For each included piece, the URMP dataset provides the musical score in MIDI format, the audio recordings of the individual tracks, the audio and video recordings of the assembled mixture, and ground-truth annotation files including frame-level and note-level transcriptions. In creating the URMP dataset, a key challenge we encountered and overcame was the synchronization of individually recorded instrumental sources of a piece while maintaining the expressiveness seen in professional chamber music performances. We present our attempts and reflections on addressing this challenge. We hope that our experience can shed some light on addressing this key issue in creating polyphonic music datasets that contain isolated source recordings. We also present some sample preliminary results from analysis on the URMP dataset. To our knowledge, the URMP dataset is the first musical performance dataset that provides ground truth audio-video recordings separately for each instrumental source. We hope that the URMP dataset will help the research community move forward along the direction of multimodal musical performance analysis.

II. Review of Musical Performance Datasets

Musical performance datasets are not easy to create because recording musical performances and annotating their ground-truth labels (e.g., pitch, chord, structure, and mood) requires musical expertise and is very time consuming. Commercial recordings generally cannot be used due to copyright issues. Recording musical performances in research labs is subject to the availability of musicians and recording facilities. Additional challenges exist in ensuring proper synchronization between different instrument parts recorded in isolation (e.g., for musical source separation). The annotation process often requires experienced musicians to listen through the musical recording multiple times. It is especially difficult when the annotations are numerical and at a temporal resolution on the order of milliseconds (e.g., pitch transcription, audio-score alignment). Therefore, musical performance datasets are scarce and their sizes are relatively small. In this section, we briefly review several commonly-used musical performance datasets that are closely related to the URMP dataset, i.e., those that can support music transcription, source separation, and audio-score alignment tasks. Most are single modality and only contain the audio. We are only aware of two audio-visual musical performance datasets, which we include here as well. A summary of these datasets is provided in Table I.

The first batch of datasets are single-track recordings with MIDI transcriptions for music transcription research. A large existing portion of music transcription datasets focus on piano music [19]–[21]. To ease the ground-truth transcription annotation process, they often have a musician perform on a MIDI keyboard and then use a Disklavier to render acoustical recordings from the MIDI performance. The MIDI performance naturally serves as the ground-truth transcription of the audio recording. There are also several multi-instrument polyphonic datasets for music transcription [22], [23]. In this case, obtaining ground-truth transcriptions is much more difficult even though MIDI scores are available, because the MIDI score has different temporal dynamics from the music performance. Manual or semi-automatic alignment between the score and the performance is required to obtain the ground-truth transcription. For the RWC dataset [22], the alignment is provided by Ewert et al. using audio-score alignment followed by manual corrections [33]. For the Su dataset [23], a professional pianist was employed to follow and play the music on an electric piano to generate well-aligned ground-truth transcriptions.

The second batch of datasets are multi-track recordings, where each instrumental source is on one track. The largest dataset is MedleyDB [24]. It contains multi-track audio recordings of 108 pieces with various styles together with the melody pitch contour and instrument activity annotations. The second largest dataset is the Structural Segmentation Multitrack Dataset (SSMD) [25], which contains multi-track audio recordings of 103 rock and pop songs, together with structural segmentation annotations. Most recordings of MedleyDB and SSMD are from third-party musical organizations (e.g., commercial or non-profit websites, recording studios). This relieved the burden of recoding by the researchers themselves. The other multi-track datasets are of a much smaller scale. MASS dataset [26] contains several raw and effects-processed multi-track audio recordings. Mixploration dataset [27] contains 3 raw multi-track audio recordings together with a number of mixing parameters. The Wood Wind Quintet (WWQ) dataset [28] contains individual recordings of 1 quintet. The first 30 seconds has been used as the development set for the MIREX[1] Multiple Fundamental Frequency Estimation and Tracking task since 2007. The TRIOS dataset [29] contains 5 multi-track recordings of musical trios together with their MIDI transcriptions. The Bach10 dataset [30] contains 10 multi-track instrumental recordings of J.S. Bach four-part chorales, together with the pitch and note transcriptions and the ground-truth audio-score alignment.

The only prior musical performance datasets that contain both audio and video recordings are the ENST-Drums dataset [31] and the Multimodal Guitar dataset [32]. The ENST-Drums dataset was created to support research on automatic drum transcription and processing. It contains mixed stereo audio tracks, audio and video recordings of each instrument of a drum kit playing different sequences including individual strokes, phrases, solo, and accompaniment parts. All instruments were recorded simultaneously using 8 microphones (1

[1] http://www.music-ir.org/mirex/wiki/MIREX_HOME

TABLE I: Summary of several commonly used musical performance datasets for music transcription, source separation, audio-score alignment, and audio-visual analysis.

Name	Type	# Pieces	Total duration	Content
MAPS [19]	Piano, single-track	270	66,880 s	Audio, MIDI transcription
LabROSA Piano [20]	Piano, single-track	130	9,600 s	Audio, MIDI transcription
Score-informed Piano Transcription [21]	Piano, single-track	7	386 s	Audio, MIDI transcription, performance error annotation
RWC (Classical & Jazz) [22]	Multi-instr., single-track	100	33,174 s	Audio, MIDI transcription
Su [23]	Multi-instr., single-track	10	301 s	Audio, MIDI transcription
MedleyDB [24]	Multi-track	108	26,831 s	Mixed and source&audio, melody F0, instrument activity
Structural Segmentation Multi-track Dataset (SSMD) [25]	Multi-track	104	24,544 s	Mixed and source audio, structural annotation
MASS [26]	Multi-track			Mixed and source audio
Mixploration [27]	Multi-track	3	73 s	Source audio, mixing parameters
Wood Wind Quintet (WWQ) [28]	Multi-track	1	54 s	Mixed and source audio
TRIOS [29]	Multi-track	5	190 s	Mixed and source audio, MIDI transcription
Bach10 [30]	Multi-track	10	330 s	Mixed and source audio, pitch and note transcription, audio-score alignment
ENST-Drums [31]	Drum kit, multi-track	N/A	13,500 s	Mixed and source audio, source video
Multimodal Guitar [32]	Guitar, single-track	10	600 s	Audio, video

for each instrument plus three extra at different locations) and 2 cameras (front and right side views). Since different instruments were not recorded in isolation, there is some leakage across different instruments. The multimodal Guitar dataset [32] contains 10 audio-visual recordings of guitar performances. The audio was recorded using a contact microphone to capture the vibration of the guitar body and attenuate the effects of room acoustics and sound radiation. The video was recorded using a high-speed camera with markers attached on joints of the player's hands and the guitar body to facilitate hand and instrument tracking.

A. Synchronization Challenges in Creating Multi-track Datasets

It is the coordination between simultaneous sound sources that differentiates music from general polyphonic acoustic scenes. One important aspect of this coordination is synchronization, which is typically accomplished in real-world musical performances by players rehearsing together prior to a performance. During the performance, players also rely on auditory and visual cues to adjust their speed to other players. For large ensembles such as a symphony orchestra, a conductor sets the synchronization.

In order to have a musical performance dataset simulate real-world scenarios, good synchronization between different instrumental parts is desired. However, creating a multi-track dataset without leakage across different tracks is challenging. This is because different instruments need to be recorded separately and players cannot rely on interactions with other players to adjust their timing. In this subsection, we review existing approaches that researchers have explored in ensuring the synchronization when recording multi-track datasets.

For SSMD [25], MASS [26], Mixploration [27], and a large portion of MelodyDB [24], recordings were obtained from professional musical organizations and recording studios instead of being recorded in a laboratory setting. The pieces are also mostly rock and pop songs, which have a steady tempo, making synchronization easier. In fact, in the music production industry, pop music is almost always produced by first recording each track in isolation and then mixing them and adding effects. This procedure, however, does not apply to classical music, which involves much less processing. Different instrumental parts of a classical music piece are almost always recorded together. This is why these datasets do not contain many classical ensemble pieces. The multi-track recordings in ENST-Drums [31] were recorded using 8 microphones simultaneously, hence no synchronization issue had to be dealt with during recording. However, leakage between microphones is inevitable, which makes the dataset less desirable for source separation research.

Existing multi-track datasets that have dealt with the synchronization issue are WWQ [28], TRIOS [29], and Bach10 [30]. WWQ [28] only has one quintet piece, and the recording process had two stages. In the first stage the performers played together with separate microphones, one for each instrument. Audio leakage inevitably existed in these recordings but they served as a basis for synchronization in the second stage. In the second stage players recorded their parts in isolation while listening to a mix of the other player's recordings in the first stage through headphones. Because these players had rehearsed together and listened to their own performance (the

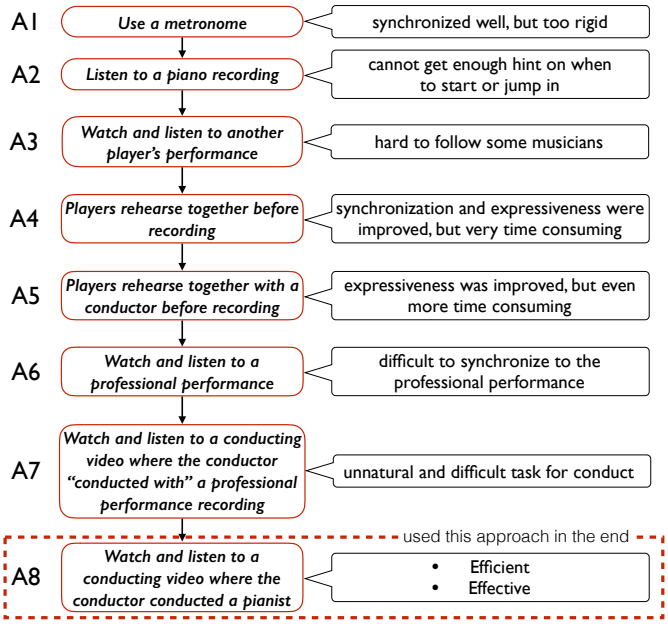
A1
Use a metronome
synchronized well, but too rigid
A2
Listen to a piano recording
cannot get enough hint on when to start or jump in
A3
Watch and listen to another player's performance
hard to follow some musicians
A4
Players rehearse together before recording
synchronization and expressiveness were improved, but very time consuming
A5
Players rehearse together with a conductor before recording
expressiveness was improved, but even more time consuming
A6
Watch and listen to a professional performance
difficult to synchronize to the professional performance
A7
Watch and listen to a conducting video where the conductor "conducted with" a professional performance recording
unnatural and difficult task for conduct
used this approach in the end
A8
Watch and listen to a conducting video where the conductor conducted a pianist
Efficient
Effective

Fig. 1: A summarized flowchart of all attempts used to solve the synchronization issue.

first-stage recordings), the synchronization among individual recordings in the second stage was very accurate. In TRIO [29], for each piece, a synthesized audio recording was first created for each instrument from the MIDI score. Each player then recorded his part in isolation while listening to the mix of the synthesized recordings of other parts synchronized with a metronome through headphones. Although the players did not rehearse together prior to the recording, the synchronization was not bad as all the pieces have a steady tempo. In Bach10 [30], instead of using the synthesized recordings and a metronome as the synchronization basis, each player listened to the mix of all previously recorded parts. The first player, however, determined the temporal dynamics and did not listen to anyone else resulting in a less-than-ideal synchronization in Bach10. Due to significant variation in the tempo, a listener could easily find many places where notes were not articulated together. In fact, each piece contains several fermata signs, where notes were prolonged beyond their normal duration when the performance was held.

III. Our Attempts to Address the Synchronization Issue

Similar to the creation of existing multi-track datasets, the creation of URMP audio-visual dataset faced the synchronization issue. This issue was even more significant because of the following seemingly conflicting goals: 1) *Efficiency*. Our goal was to create a large dataset containing dozens of pieces with different instrument combinations. We also hoped that each player could participate in the recording of multiple pieces. Therefore, it would be too difficult and time consuming to arrange players to rehearse together before the recording for each piece, which was the approach adopted by the creation of WWQ [28]. 2) *Quality*. We wanted the players to be as expressive as what they would be in real musical concerts. This required them to vary the tempo and dynamics significantly throughout a piece. However, without the live interaction between players, this goal made the synchronization more difficult. We tried different ways to overcome this challenge and eventually found a way that achieved both good efficiency and quality. We present our attempts here and hope that this will give some insights into the dataset creation problem. Figure 1 summarizes our attempts.

A1) The first approach that we tried was to pre-generate a beat sequence using an electronic metronome, and then have each player listen to the beat sequence through a headphone while recording his/her part. Different instrumental tracks were thus synchronized through the common beat sequence. We tested this approach on a violin-cello duet (Minuet in G major by J. S. Bach). Although the synchronization was good, we found that the performance was too rigid and did not reach our desired level of expressiveness.

A2) In order to have better expressiveness, we replaced the beat sequence with a pianist's recording of the piece. The pianist played both the violin and cello parts simultaneously on an electric keyboard and audio was rendered using a synthesizer. The pianist varied the tempo and dynamics throughout the piece to increase the expressiveness. However, when the players listened to the piano recording and recorded their individual parts, they reported that it was quite difficult to synchronize to the piano. They did not get enough hint on when to start in the beginning nor when to jump in after a long break.

A3) Another way we tried was to have each player to synchronize to another player's performance directly during the recording, similar to Bach10's approach [30]. However, each player watched another player's video instead of just listening to the audio. We again tested on the minuet piece where the cello was recorded first followed by the violin. We found that watching the video helped to improve the synchronization, mainly at the beginning of the piece. However, the violin player still reported that it was hard to follow the cello player's performance and the synchronization was not satisfactory.

A4) In this attempt we tried the WWQ's approach [28] and chose another violin-cello piece (Melody by Schumann). It turned out that both synchronization and expressiveness were greatly improved.

A5) Improving upon A4, we further added a conductor to the rehearsal process to improve the expressiveness. The conductor varied the tempo and dynamics significantly, and also vocalized several beats before the start. The vocalization was also recorded during the rehearsal, and it gave players a better hint on when to start. It turned out that both the synchronization and expressiveness were great. However, this approach was very time consuming and would not scale to a large dataset which we aimed at. Nevertheless, these attempts gave us an idea about the highest level of synchronization and expressiveness we can expect in creating such datasets.

A6) Attempts from A1-A5 reveals the advantage of using visual cues along with the auditory cues. Building on this idea, we asked each player to watch and listen to all parts of the real video performance downloaded from YouTube during recording. We tested this approach on a string quartet (Art of

Fugue No. 1 by J. S. Bach). The players' experience revealed that, due to the lack of clarity in the visual cues, it was difficult to synchronize to the professional video even after repeatedly watching the video.

A7) This attempt focused on shifting the burden of synchronization from players by using a professional conductor to conduct the YouTube video. The conducting video along with the played-back audio was recorded and each player was asked to watch the conducting video during recording. Although this approach helped players synchronize to professional music performance, it made the conductor's job very difficult. The conductor had to memorize the temporal dynamics of the performance and behave in a timely fashion following the performance. This was very non-intuitive as conductors should conduct a performance instead of "being conducted" by a performance.

A8) In order to strike a balance between the conductor's and the players burdens, our final attempt had two key steps. In the first step, we asked the conductor to conduct a pianist performing the piece, and recorded the conductor's video and the piano audio with a digital camera. The conductor varied the tempo and dynamics and the pianist adjusted the performance accordingly. The conductor also gave cues to different instrumental parts to help players jump in after a long rest. As a second step, we asked each player to watch the conducting video on a laptop and listen to the audio through a headphone during the recording. This process turned out to be both efficient and effective. The conductor, pianist, and instrumental players all felt that this was the best approach.

IV. The Dataset

This section explains in detail the execution of the two key steps introduced in A8 in Section III. It covers the entire process of dataset creation, from piece selection and musician recruitment, to recording, post-production, and ground-truth annotation. The whole process is summarized in Fig. 2, taking a duet as an example.

A. Piece Selection

Since we wanted to have a good coverage of polyphony, composers, and instrument combinations in our dataset, the choice of pieces to record was the first question we needed to answer. We wanted to choose such pieces that are (a) not too simple, which would be boring (b) not too difficult, which would require too much practice for the musicians. Regarding the duration, we thought 1 to 2 minutes would be ideal. Bearing these guidelines in mind, we searched for sheet music resources and found a great sheet music website[2], which provides thousands of simplified and rearranged musical scores of different polyphony, styles, composers, and instrument combinations. We selected a number of classical ensemble pieces, covering duets, trios, quartets, and quintets. The duration of these pieces ranges from 40 seconds to 4.5 minutes, and most are around 2 minutes. They cover different instrument combinations, including string groups, woodwind groups, brass groups, and mixed groups. To create more instrument combinations, we further varied the instruments in several pieces. This resulted in 44 recorded pieces, as listed in Table II.

[2]http://www.8notes.com

B. Recruiting Musicians

Creating the URMP dataset required three kinds of musicians: a conductor, a pianist, and musicians who recorded the instrumental parts of the pieces. All of the musicians were either students from the Eastman School of Music or members of various music ensembles and orchestras. The conductor had more than 20 years of conducting experience. The pianist is a graduate student majoring in piano performance. Background statistics of the instrumental players are summarized in Table III. All of the musicians signed a consent form and received a small monetary compensation for their participation.

TABLE III: Demographic statistics of musicians who recorded instrumental parts of the URMP dataset.

Current position	
Undergraduate	19
MS student	3
PhD student	1
Professional vs. Nonprofessional	
Professionals	17
Nonprofessionals	6
Playing experience	
>20 years	2
10-20 years	9
Unknown	12
Total Number of Musicians	23

C. Recording Conducting Videos

For each piece, a video consisting only of a conductor conducting a pianist playing on a Yamaha grand piano was recorded to serve as the basis for the synchronization of different instrumental parts. These conducting videos were recorded in a 25'x18' recording studio using a Nikon D5300 camera and its embedded microphone. Before recording each piece, the conductor and the pianist rehearsed several times and the conductor always started with several extra beats for the pianist and the other musicians to follow. The tempo of each piece was set by the conductor and the pianist together after considering the tempo notated in the score. All repeats within a piece were reserved for integrity. Grace notes, fermatas, and tempo changes notated in the score were also implemented for high expressiveness.

D. Recording Instrumental Parts

The recording of the isolated instrumental parts was conducted in an anechoic sound booth with the floor plan as shown in Figure 3. The wall behind the player was covered by a blue curtain and the lighting of the sound booth was through fluorescent lights affixed at wall-ceiling intersections around the room. We placed a Nikon D5300 camera on the front-right side of the player to record the video with a 1080p resolution. Since the built-in microphone of the camera does

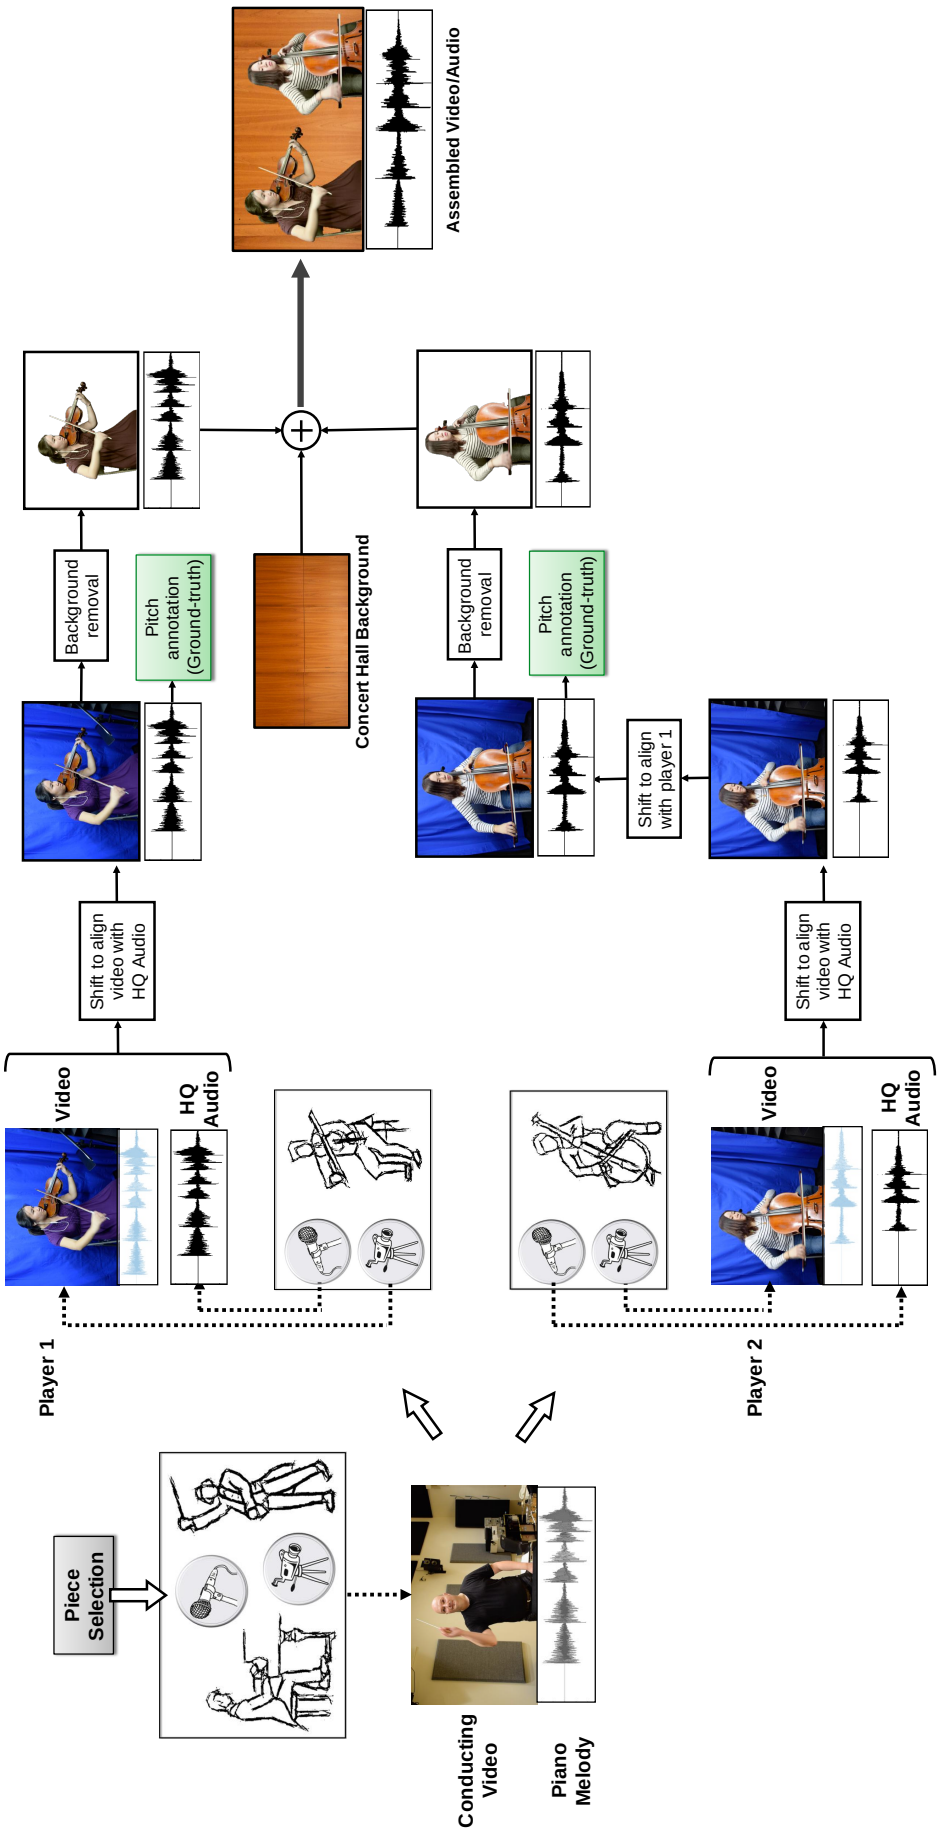

Fig. 2: The general process of creating one piece (a duet in this example) of the URMP dataset.

TABLE II: Musical pieces and their instrumentation contained in the URMP dataset.

Index	Composer and title	Instruments	Length (min:sec)
	Duet (11 videos)		
1	Holst - Jupiter from the Planets	Violin, Cello	1:03
2	Mozart - Theme from sonata k.331	Violin I, Violin II	0:46
3	Tchaikovsky - Dance of the Sugar Plum Fairy	Flute, Clarinet	1:37
4	Beethoven - Allegro for a Musical Clock	Flute I, Flute II	2:37
5	Scott Joplin - The Entertainer	Trumpet I, Trumpet I	1:27
6	Scott Joplin - The Entertainer	Alto Sax I, Alto Sax II	1:28
7	Bach - Air on a G String	Trumpet, Trombone	4:30
8	Vivaldi - Spring from The Four Seasons	Flute, Violin	0:35
9	Bach - Jesu Joy of Mans Desiring	Trumpet, Violin	3:19
10	Handel - March from Occasional Oratorio	Trumpet, Saxophone	1:43
11	Schubert - Ave Maria	Oboe, Cello	1:44
	Trio (12 videos)		
12	Vivaldi - Spring(longer version)	Violin I, Violin II, Cello	2:12
13	Mendelssohn - Hark the Herald Angels Sing	Violin I, Violin II, Viola	0:47
14	Tchaikovsky - Waltz from Sleeping Beauty	Flute I, Flute II, Clarinet	1:34
15	Haydn - Theme	Trumpet I, Trumpet II, Trombone	0:54
16	Haydn - Theme	Trumpet I, Trumpet II, Alto Sax	0:54
17	Mendelssohn - Nocturne	Flute, Clarinet, Violin	1:36
18	Mendelssohn - Nocturne	Flute, Trumpet, Violin	1:36
19	Fauré Pavanne	Clarinet, Violin, Cello	2:13
20	Fauré Pavanne	Trumpet, Violin, Cello	2:13
21	Handel - La Rejouissance	Clarinet, Trombone, Tuba	1:16
22	Handel - La Rejouissance	Soprano Sax, Trombone, Tuba	1:16
23	Handel - La Rejouissance	Clarinet, Alto Sax, Tuba	1:16
	Quartet (14 videos)		
24	David Bruce - Pirates	Violin I, Violin II, Viola, Cello	0:50
25	David Bruce - Pirates	Violin I, Violin II, Viola, Alto Sax	0:50
26	Grieg - In the Hall of the Mountain King	Violin I, Violin II, Viola, Cello	1:25
27	Grieg - In the Hall of the Mountain King	Violin I, Violin II, Viola, Alto Sax	1:25
28	Bach -The Art of Fugue	Flute, Oboe, Clarinet, Bassoon	2:55
29	Bach - The Art of Fugue	Flute I, Flute II, Oboe, Clarinet	2:55
30	Bach - The Art of Fugue	Flute I, Flute II, Oboe, Soprano Sax	2:55
31	Dvořák - Slavonic Dance	Trumpet I, Trumpet II, Horn, Trombone	1:22
32	Bach -The Art of Fugue	Violin I, Violin II, Viola, Cello	2:54
33	Beethoven - Fur Elise	Horn, Trumpet I, Trumpet II, Trombone	0:48
34	Bach - The Art of Fugue	Trumpet I, Trumpet II, Horn, Trombone	2:54
35	Purcell - Rondeau from Abdelazer	Violin I, Violin II, Viola, Bass	2:08
36	Purcell - Rondeau from Abdelazer	Violin I, Violin II, Viola, Cello	2:08
37	Purcell - Rondeau from Abdelazer	Flute, Clarinet, Viola, Violin	2:08
	Quintet (7 videos)		
38	Parry Jerusalem	Violin I, Violin II, Viola, Cello, Bass	1:59
39	Parry Jerusalem	Violin I, Violin II, Viola, Sax, Bass	1:59
40	Allegri - Miserere Mei Deus	Flute I, Flute II, Clarinet, Oboe, Bassoon	0:40
41	Allegri - Miserere Mei Deus	Flute I, Flute II, Soprano Sax, Oboe, Bassoon	0:40
42	Bach - Arioso from Cantata BWV 156	Trumpet I, Trumpet II, Horn, Trombone, Tuba	1:40
43	Saint-Sans - Chorale	Trumpet I, Trumpet II, Horn, Trombone, Tuba	2:40
44	Mozart - String Quintet K515 2nd Mvt	Violin I, Violin II, Viola I, Viola II, Cello	3:40

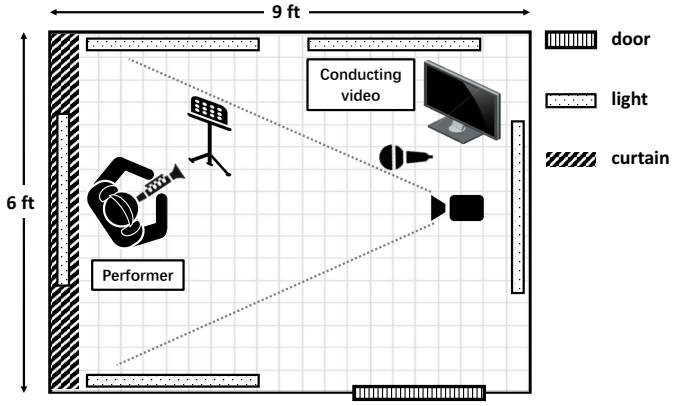

Fig. 3: The floor plan of the sound booth (top-down view).

not achieve adequate audio quality, we also used an Audio-Technical AT2020 condenser microphone to record high-quality audio with a sampling rate of 48 KHz and a bit depth of 24. We connected the microphone to a laptop computer running Audacity™ [34] for the audio recording thereby making camera and the stand-alone microphone independently controllable.

During the recording, the player watched the conducting video on a laptop with a 13-inch screen placed about 5 feet in front of the player. The player also listened to the audio track of the conducting video through a blue-tooth wireless earphone. All musicians reported that watching and listening to the conducting videos greatly helped them synchronize their performance with the piano. For simpler pieces, the recording was finished in one shot; while for long and difficult pieces, several shots were conducted before we approved the quality of the recordings.

E. Mixing and Assembling Individual Recordings

For each instrumental part, we first replaced the low-quality audio in the original video recording with the high-quality audio recorded using the stand-alone microphone. Since the camera and the stand-alone microphone were controlled independently, the video and the high-quality audio recordings did not start at the same time and should be temporally aligned. We used the "synchronize clips" function of the Final Cut Pro software [35] to find the alignment and then did the replacement.

We then assembled individual instrumental recordings. Although the individual instrumental parts of each piece were all aligned with the conducting video, the conducting video started to play at different times relative to the beginning of the individual recordings. We need to temporally shift the individual recordings to align them. We manually did so using the Final Cut Pro software. The fast sections with clear note onsets were especially helpful for us to make the shift decisions. We also manually adjusted the loudness of some players of some pieces to achieve a better balance. We prefer this subjective adjustment instead of an objective normalization of their root-mean-square values to make the balance of different parts more natural. Finally, we assembled the synchronized individual video recordings into a single ensemble recording. In the video, all players were arranged at the same level from left to right. The order of the players followed the order of score tracks. The audio is the mixture (addition) of the individual high-quality audio recordings.

F. Video Background Replacement

In order to make the assembled videos look more natural and similar to live ensemble performances, we used chroma keying [36] to replace the blue curtain background with a real concert hall image.

We use the Final Cut Pro software package for video compositing purpose. The blue background in the videos is unevenly lit and has significant textural variation. To overcome this, we did color correction as a pre-processing step followed by chroma keying. By adjusting the keying and color correction parameters and by setting suitable selection spatial masks, we were able to get a good separation between foreground and background. Once the foreground was extracted, we used more realistic image as the background for the composite video. The background photo was captured from Hatch Recital Hall[3] using a Nikon D5300 camera.

G. Ground-truth Annotation

We provide the musical scores for all the 44 pieces. As shown in Table II, some pieces are derived from the same original piece with different instrument arrangements and/or keys.

We also provide ground-truth pitch annotation for each audio track. This annotation process was performed on each single audio track using the Tony melody transcription software [37], which implements pYIN [38], a state-of-the-art frame-wise monophonic fundamental frequency (F0) estimation algorithm. For each audio track, we generated two files: a frame-level pitch trajectory and a note sequence. The pitch trajectory was first calculated with a frame length of 46 ms and a hop size of 5.8 ms, and then interpolated to change the hop size to 10 ms according to the standard format of ground-truth pitch trajectories in MIREX. The note sequence was extracted by the Tony software by using Viterbi decoding of a hidden Markov model [39]. The pitch of each note takes un-quantized frequencies. Many manual corrections were then performed on both the frame-level and note-level annotations to guarantee the quality of the ground-truth data.

H. Dataset Content

The URMP dataset contains audio-visual recordings and ground-truth annotations for all the 44 pieces, each of which is organized in a folder with the following contents:

- **Score:** we provide both sheet score in PDF format and MIDI score. The sheet score is directly generated from the MIDI score using Sibelius 7.5 [40] with minor adjustment (clef, key set, note spelling, etc.) for display purpose only. The order of score tracks is arranged musicologically;

[3]http://www.esm.rochester.edu/concerts/halls/hatch/

- **Audio:** individual and mixed high-quality audio recordings in WAV format, with a sampling rate of 48 KHz and a bit depth of 24. The naming convention of individual tracks follow the same order as the tracks in the score.
- **Video:** assembled video recordings in MP4 format with H264 codec. Video files have the resolution of 1080P (1920×1080), and frame rate of 29.97fps. Players are rendered horizontally, from left to right, following the same order as the tracks in the score.
- **Annotation:** ground-truth frame-level pitch trajectories and note-level transcriptions of individual tracks in ASCII delimited text format.

A sample piece of the dataset is available at online[4]. The full dataset, which has a size of 12.5 GB, will be available for download when the paper is published.

I. Evaluating the Synchronization Quality of Dataset

Because maintaining the synchronization among different instruments is the main challenge in creating the URMP dataset, we conducted subjective evaluations to compare the synchronization quality of this dataset with that of Bach10 [30] and WWQ [28]. Both datasets have been used in the development and/or testing phase for the MIREX Multi-F0 Estimation & Tracking task in the past. We recruited 8 subjects who were students at the University of Rochester. None of them were familiar with these datasets. For each subject, we randomly chose pieces from these datasets to form 4 triplets, one piece from each dataset. We then asked the subjects to listen to the three pieces of each triplet and rank their synchronization quality. Table IV shows the ranking statistics. It can be seen that out of the 32 rankings (4 rankings per subject for 8 subjects), URMP ranked first 9 times, and ranked second 17 times. The process we developed for recording the URMP dataset achieves a much higher synchronization quality than the Bach10 pieces. While the WWQ dataset was rated better, it contains only one piece that was recorded after players rehearsed together in advance, a methodology that does not scale to a larger dataset.

TABLE IV: Subjective ranking results of the synchronization quality of three datasets provided by eight subjects.

Rank	#1	#2	#3
URMP	9	17	6
WWQ [28]	22	9	1
Bach10 [30]	1	6	25

V. Applications of The Dataset

As the first audio-visual polyphonic music performance dataset with individually separated sound sources, the URMP dataset not only serves for the evaluation of approaches to many existing research tasks in Music Information Retrieval (MIR) such as source separation and music transcription, but it also encourages the transformation of many existing problems and the creation of some fundamentally new approaches via the use of the visual modality of the musical performances. In this section, we present our preliminary studies on some of these research problems, hoping this will prompt further discussion and interactions in the research community.

[4]http://www.ece.rochester.edu/projects/air/projects/datasetproject.html

A. Preliminary Analyses on Playing Styles

We used optical flow techniques to track the bowing hand motions of string instrument players. Figure 4(b) shows an the motion vectors estimated from the highlighted region in Figure 4(a). Figure 4(c) shows the distribution of all global motion vectors across all time frames, where each dot (i.e., a global motion vector) represents the average motion vector for all the pixels in one frame. The green line is the main motion direction obtained by using Principal Component Analysis [41].

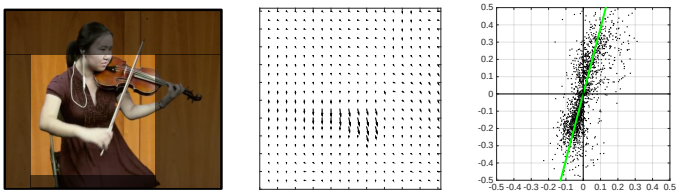

(a) (b) (c)

Fig. 4: Example of optical flow analysis. (a) A video frame with a highlighted region of interest. (b) Vector field representation of the estimated motion vectors in the highlighted region. (c) Aggregated global motion vectors of all frames with the green line indicating the main motion direction.

Figure 5 shows this motion feature for two violin players in three different musical pieces, where they played different parts. It can be seen that the motion vectors of the first player exhibit a dominant principal direction, while those of the second player are more randomly distributed. Although this difference may be due to the differences of their instrumental parts, we argue that it actually mainly reflects their unique playing habits, which can be related to their professionalism. The same technique can be applied to analyze the bowing & plucking playing techniques. Figure 6 shows the global motion vectors for the same player with different playing techniques. They are very different, in terms of both the scale and the principal directions.

B. Audio-visual Source Association

Based on the above mentioned methods, we have addressed the source association problem in video recordings of multi-instrument music performances by synergistic combination of analysis of the audio and video modalities [17]. The framework is shown in Figure 7. The key idea in this work lies in recognizing that characteristics of sound sources corresponding to individual instruments exhibit strong temporal interdependence with motion observed on these instrument players. Experiments were performed on 19 musical pieces from the dataset that involve at most one non-string player. Results show that 89.2% of individual tracks are correctly associated. This approach enables novel music enjoyment experiences by allowing users to target an audio source by clicking on the player in the video to separate/enhance it [42].

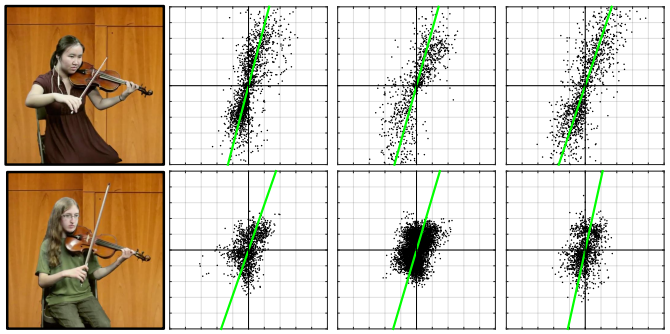

(a) (b)

Fig. 5: Global motion vectors of two violinists playing different parts of three musical pieces.

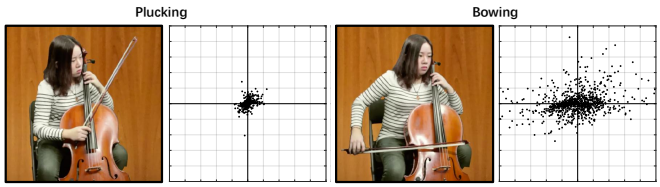

(a) (b)

Fig. 6: Distributions of the global motion vectors for the same player with (a) plucking and (b) bowing techniques.

C. Audio-visual Multi-pitch Analysis

In another application, we performed multi-pitch analysis on string ensembles by leveraging the visual information [8]. The central idea in this work is to determine the Play/Non-play (P/NP) activity in each video frame and utilize the P/NP labels to obtain the information about instantaneous polyphony and active sources, which is provided as input to audio-based Multi-pitch Estimation (MPE) and Multi-pitch Streaming (MPS) modules, respectively, as shown in Figure 8. Experiments conducted on 11 string ensembles reveal us an overall accuracy of 85.3% on video-based P/NP activity detection, and a statistically significant improvement in the MPE and MPS accuracy at a significance level of 0.01 in most cases.

VI. Conclusion

In this paper, we presented the URMP dataset, an audio-visual musical performance dataset that is useful for a broad range of research topics including source separation, music transcription, audio-score alignment, and musical performance analysis. In particular, we discussed the main challenge and our different attempts for ensuring the synchronization of individual instrumental parts that were separately recorded. The final and successful approach we adopted involved having individual instrument players watch and listen to a pre-recorded conducting video when recording their individual parts. We also briefly described our preliminary work on some music information retrieval (MIR) tasks using the URMP dataset. We anticipate that the URMP dataset will become a valuable resource for researchers in the field of music information retrieval and multimedia.

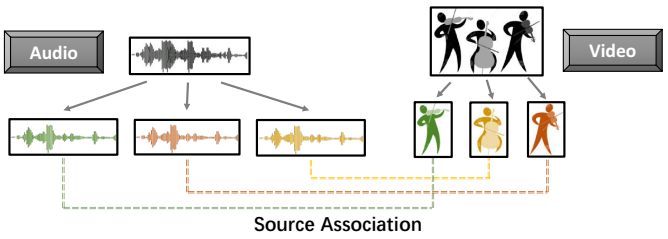

Fig. 7: Illustration of the audio-visual source association system.

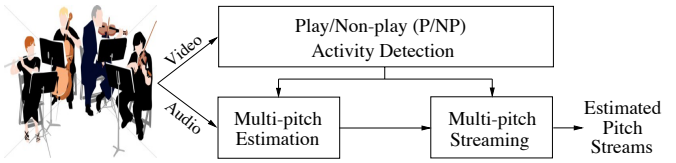

Fig. 8: Framework of the audio-visual multi-pitch analysis system.

Acknowledgment

We would like to thank Sarah Rose Smith for playing the cello, and graciously participating in our various attempts at cross-player synchronization. We also would like to thank Andrea Cogliati for recording the conducting videos and Dr. Ming-Lun Lee and Yunn-Shan Ma for supporting the research.

References

[1] C. Smith, http://expandedramblings.com/index.php/youtube-statistics/4/.
[2] F. Platz and R. Kopiez, "When the eye listens: A meta-analysis of how audio-visual presentation enhances the appreciation of music performance," Music Perception: An Interdisciplinary Jour., vol. 30, no. 1, pp. 71–83, 2012.
[3] C.-J. Tsay, "Sight over sound in the judgment of music performance," National Academy of Sciences, vol. 110, no. 36, pp. 14 580–14 585, 2013.
[4] B. Zhang, J. Zhu, Y. Wang, and W. K. Leow, "Visual analysis of fingering for pedagogical violin transcription," in Proc. ACM Intl. Conf. on Multimedia, 2007.
[5] M. Akbari and H. Cheng, "Real-time piano music transcription based on computer vision," IEEE Trans. on Multimedia, vol. 17, no. 12, pp. 2113–2121, 2015.
[6] M. Paleari, B. Huet, A. Schutz, and D. Slock, "A multimodal approach to music transcription," in Proc. Intl. Conf. on Image Process. (ICIP), 2008.
[7] O. Gillet and G. Richard, "Automatic transcription of drum sequences using audiovisual features," in Proc. Intl. Conf. on Acoustics, Speech, and Sig. Process. (ICASSP). IEEE, 2005.
[8] K. Dinesh, B. Li, X. Liu, Z. Duan, and G. Sharma, "Enhancing multi-pitch estimation of chamber music performances via video-based activity detection," in Proc. Intl. Conf. on Acoustics, Speech, and Sig. Process. (ICASSP). IEEE, 2017, accepted.
[9] A. Bazzica, C. C. Liem, and A. Hanjalic, "Exploiting instrument-wise playing/non-playing labels for score synchronization of symphonic music." in Proc. Intl. Soc. for Music Info. Retrieval (ISMIR), 2014.
[10] A.-M. Burns and M. M. Wanderley, "Visual methods for the retrieval of guitarist fingering," in Proc. Intl. Conf. on New Interfaces for Musical Expression (NIME), 2006.
[11] C. Kerdvibulvech and H. Saito, "Vision-based guitarist fingering tracking using a bayesian classifier and particle filters," in Advances in Image and Video Tech. Springer, 2007, pp. 625–638.

[12] D. Radicioni, L. Anselma, and V. Lombardo, "A segmentation-based prototype to compute string instruments fingering," in Proc. Conf. on Interdisciplinary Musicology (CIM), 2004.

[13] J. Scarr and R. Green, "Retrieval of guitarist fingering information using computer vision," in Proc. Intl. Conf. on Image and Vision Computing New Zealand (IVCNZ), 2010.

[14] D. Gorodnichy and A. Yogeswaran, "Detection and tracking of pianist hands and fingers," in Proc. Canadian Conf. on Comp. and Robot Vision, 2006.

[15] A. Oka and M. Hashimoto, "Marker-less piano fingering recognition using sequential depth images," in Proc. Korea-Japan Joint Workshop on Frontiers of Comp. Vision (FCV), 2013.

[16] D. Murphy, "Tracking a conductor's baton," in Proc. Danish Conf. on Pattern Recognition and Image Analysis, 2003.

[17] B. Li, K. Dinesh, Z. Duan, and G. Sharma, "See and listen: Score-informed association of sound tracks to players in chamber music performance videos," in Proc. Intl. Conf. on Acoustics, Speech, and Sig. Process. (ICASSP). IEEE, 2017, accepted.

[18] B. Yao, J. Ma, and L. Fei-Fei, "Discovering object functionality," in Proc. Intl. Conf. on Comp. Vision (CVPR), 2013.

[19] V. Emiya, R. Badeau, and B. David, "Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle," IEEE/ACM Trans. Audio, Speech, Language Process., vol. 18, no. 6, pp. 1643–1654, 2010.

[20] G. E. Poliner and D. P. Ellis, "A discriminative model for polyphonic piano transcription," EURASIP Jour. on Applied Sig. Process., vol. 2007, no. 1, pp. 154–154, 2007.

[21] E. Benetos, A. Klapuri, and S. Dixon, "Score-informed transcription for automatic piano tutoring," in Proc. European Sig. Process. Conf. (EUSIPCO). IEEE, 2012, pp. 2153–2157.

[22] M. Goto, H. Hashiguchi, T. Nishimura, and R. Oka, "RWC music database: Popular, classical and jazz music databases." in Proc. Intl. Soc. for Music Info. Retrieval (ISMIR), vol. 2, 2002, pp. 287–288.

[23] L. Su and Y.-H. Yang, "Escaping from the abyss of manual annotation: New methodology of building polyphonic datasets for automatic music transcription," in Proc. Intl. Symp. on Computer Music Multidisciplinary Research, 2015.

[24] R. M. Bittner, J. Salamon, M. Tierney, M. Mauch, C. Cannam, and J. P. Bello, "Medleydb: A multitrack dataset for annotation-intensive MIR research." in Proc. Intl. Soc. for Music Info. Retrieval (ISMIR), 2014, pp. 155–160.

[25] S. Hargreaves, A. Klapuri, and M. Sandler, "Structural segmentation of multitrack audio," IEEE/ACM Trans. Audio, Speech, Language Process., vol. 20, no. 10, pp. 2637–2647, 2012.

[26] "Mass: Music audio signal separation," http://mtg.upf.edu/download/datasets/mass.

[27] M. Cartwright, B. Pardo, and J. Reiss, "MIXPLORATION: Rethinking the audio mixer interface," in Proc. Intl. Conf. on Intelligent User Interfaces. ACM, 2014, pp. 365–370.

[28] M. Bay, A. F. Ehmann, and J. S. Downie, "Evaluation of multiple-F0 estimation and tracking systems." in Proc. Intl. Soc. for Music Info. Retrieval (ISMIR), 2009, pp. 315–320.

[29] J. Fritsch and M. D. Plumbley, "Score informed audio source separation using constrained nonnegative matrix factorization and score synthesis," in Proc. Intl. Conf. on Acoustics, Speech and Sig. Process. (ICASSP). IEEE, 2013, pp. 888–891.

[30] Z. Duan, B. Pardo, and C. Zhang, "Multiple fundamental frequency estimation by modeling spectral peaks and non-peak regions," IEEE/ACM Trans. Audio, Speech, Language Process., vol. 18, no. 8, pp. 2121–2133, 2010.

[31] O. Gillet and G. Richard, "Enst-drums: an extensive audio-visual database for drum signals processing." in Proc. Intl. Soc. for Music Info. Retrieval (ISMIR), 2006, pp. 156–159.

[32] A. Perez-Carrillo, J.-L. Arcos, and M. Wanderley, "Estimation of guitar fingering and plucking controls based on multimodal analysis of motion, audio and musical score." in Proc. Intl. Symposium on Computer Music Multidisciplinary Research (CMMR), 2015, pp. 71–87.

[33] S. Ewert, M. Müller, and P. Grosche, "High resolution audio synchronization using chroma onset features," in Proc. Intl. Conf. on Acoustics, Speech and Sig. Process (ICASSP). IEEE, 2009, pp. 1869–1872.

[34] "Audacity," http://www.audacityteam.org, accessed: Aug, 2015.

[35] "Final cut pro," http://www.apple.com/final-cut-pro/, accessed: Jul, 2016.

[36] S. Shimoda, M. Hayashi, and Y. Kanatsugu, "New chroma-key imagining technique with hi-vision background," IEEE Trans. on Broadcasting, vol. 35, no. 4, pp. 357–361, Dec 1989.

[37] M. Mauch, C. Cannam, R. Bittner, G. Fazekas, J. Salamon, J. Dai, J. Bello, and S. Dixon, "Computer-aided melody note transcription using the tony software: Accuracy and efficiency," in Proc. Intl. Conf. on Tech. for Music Notation and Representation, 2015.

[38] M. Mauch and S. Dixon, "pYIN: A fundamental frequency estimator using probabilistic threshold distributions," in Proc. Intl. Conf. on Acoustics, Speech, and Sig. Process. (ICASSP). IEEE, 2014.

[39] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," Proc. IEEE, vol. 77, no. 2, pp. 257–286, 1989.

[40] "Sibelius," http://www.avid.com/en/sibelius, accessed: Sept, 2014.

[41] I. Jolliffe, Principal component analysis. Wiley Online Library, 2002.

[42] B. Li, Z. Duan, and G. Sharma, "Associating players to sound sources in musical performance videos," in Late Breaking Demo, Intl. Soc. for Music Info. Retrieval (ISMIR), 2017.